High Resolution Climate Simulation and Regional Water Supplies

Bill Bosl, Art Mirin, Phil Duffy

Lawrence Livermore National Lab

Climate and Carbon Cycle Modeling Group

Center for Applied Scientific Computing

April 24, 2003

High-Performance Computing for Climate Modeling as a Planning Tool

- **GLOBAL WARMING IS HERE!! … so now what?**
- **How will climate change really affect societies?**
- **Effects of global climate change are local**
- **Some effects of climate change can be mitigated**
- **Requires accurate information**

Climate simulation can be used as a planning and policy tool

California Water Supply: Critical to California Economy

- Central valley: the nation's vegetable garden;
- Fifth largest economy on earth;
- Population continues to grow
- Plenty of water – distribution in time and space is the issue

Global Climate Change and Water Supply

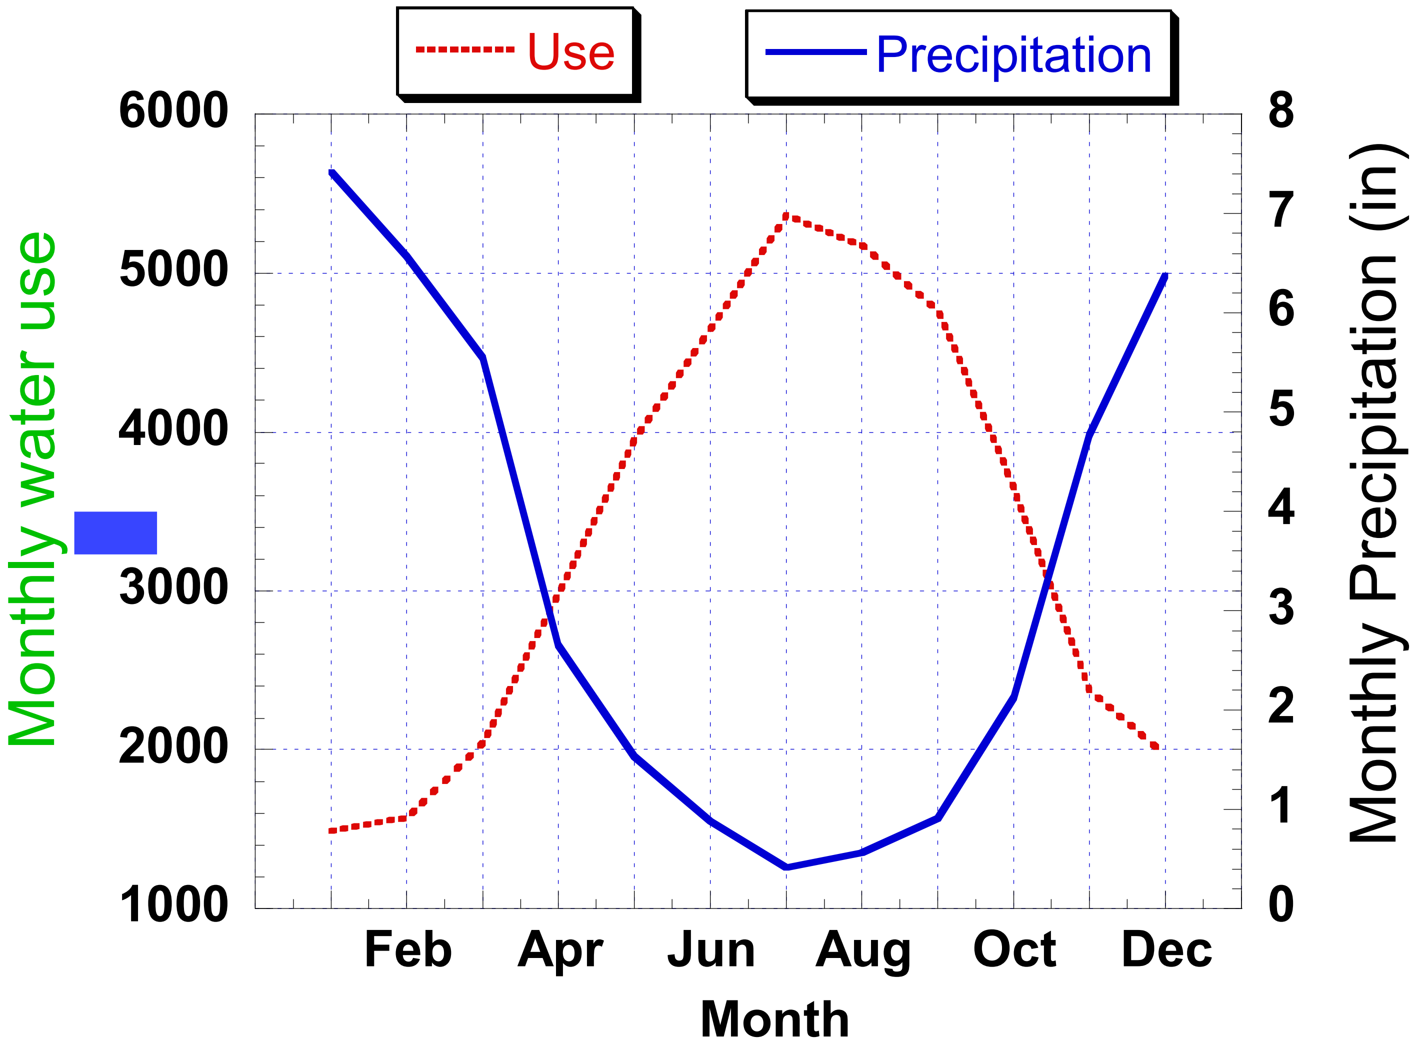

Global Climate Change: California

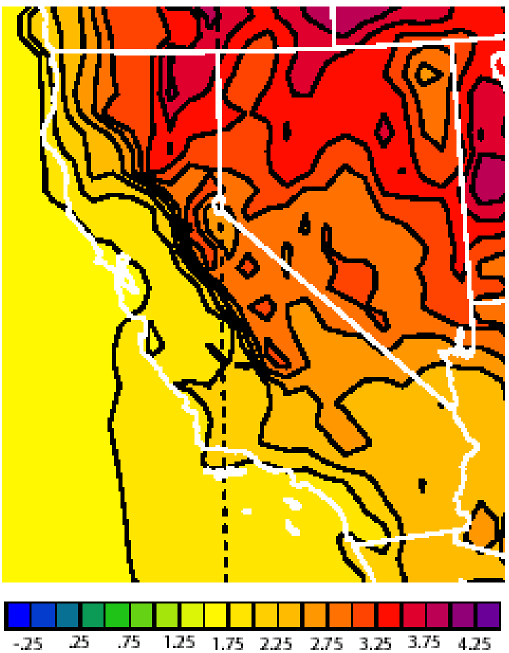

Higher Temp

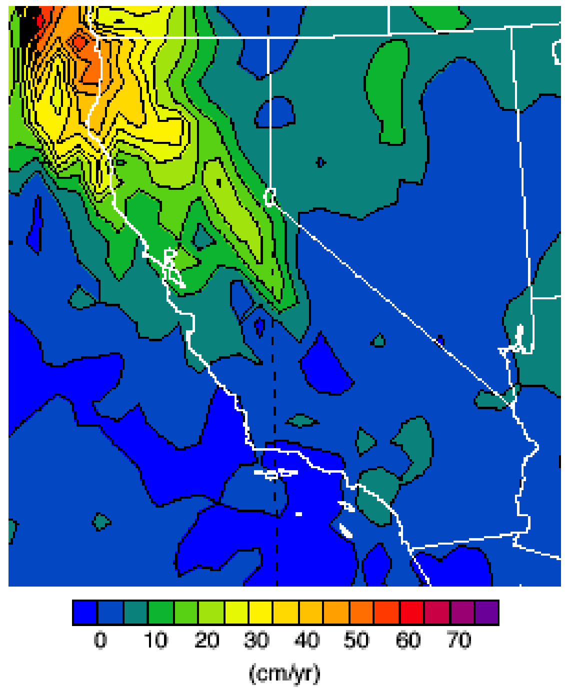

More precip

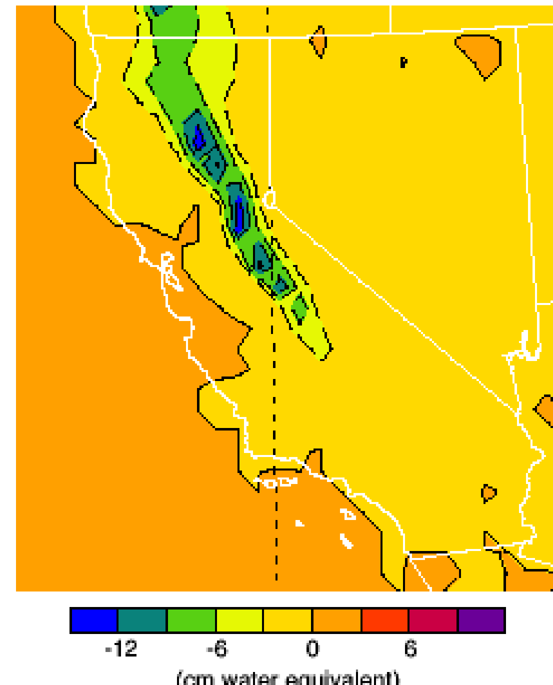

Sierra Snowpack smaller by factor of 2-3

Figures show projected changes in California region ($2xCO_2$ - $1xCO_2$) from a joint UCSC/LLNL study.

California Water Supply: Snowpack is a Huge Reservoir

CA will need >25% new reservoir capacity

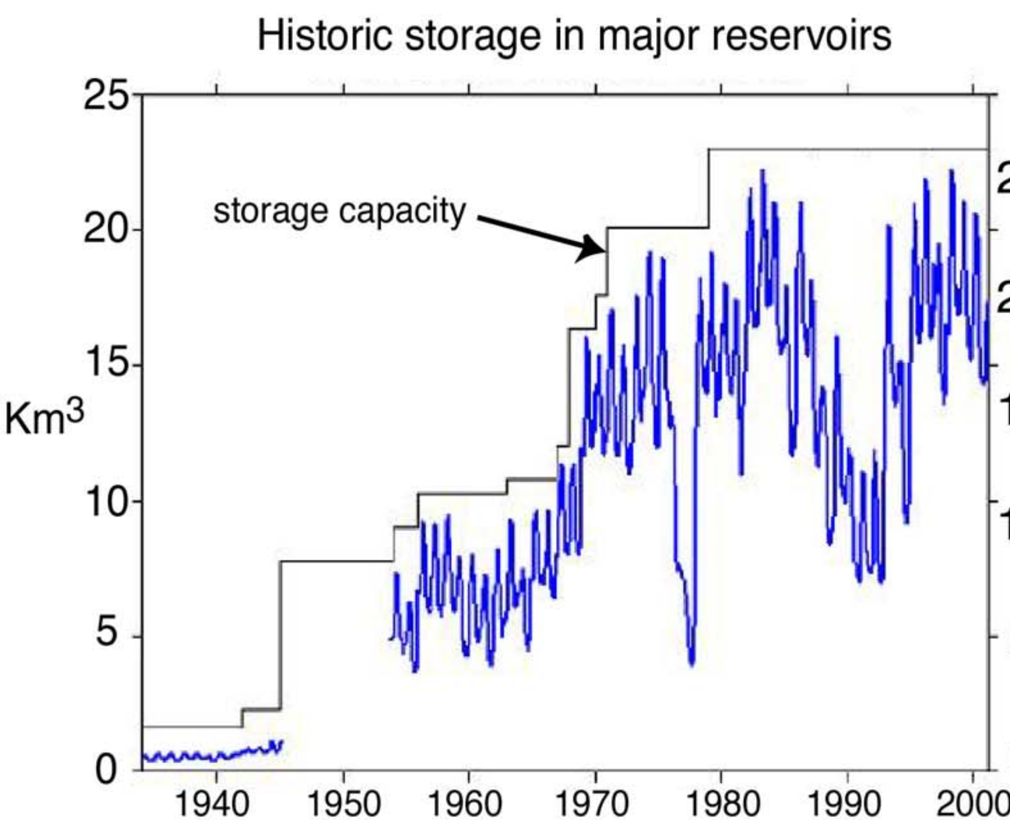

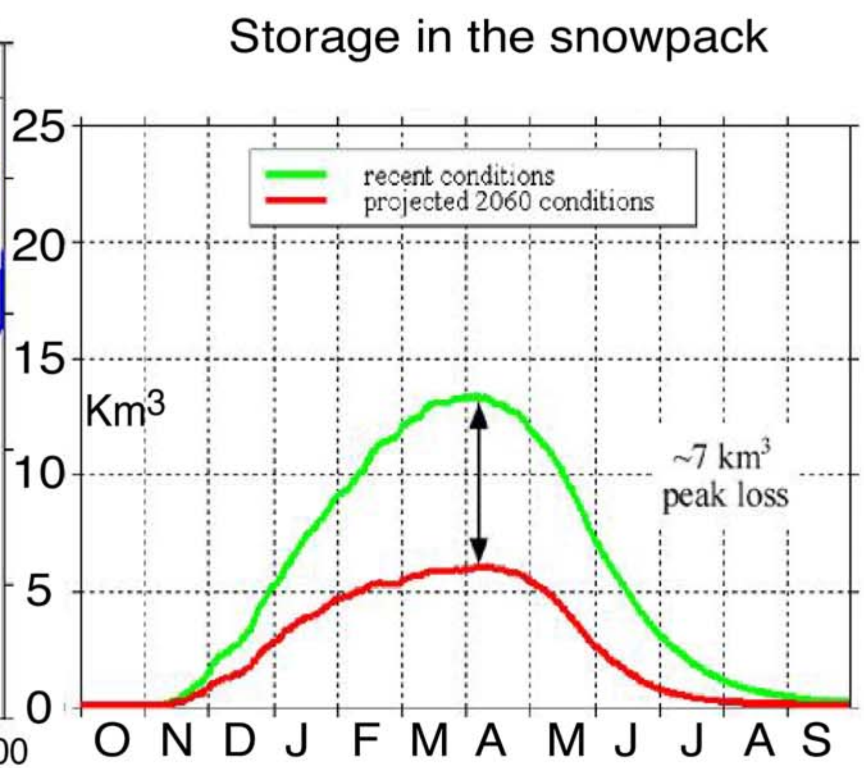

What specific questions will we address?

- **Amounts and spatial pattern of precipitation;**
- **The partitioning of precipitation between rain and snow;**
- **The water content of the Sierra snowpack;**
- **Rates and timing of river flows**

We will predict changes in peak values as well as in time-averages

New Univ of California Institute at LLNL

- Institute for Research on Climate change and its Societal Impacts (IRCCSI) will
  - Address problems of societal impacts of climate change (on agriculture, human health,water, etc.);
  - Combine high-power climate modeling at UC labs with impacts expertise at UC campuses;
  - Increase flow of UC students, postdocs, and faculty through LLNL;
- Recently approved by the UC President’s Office.

We will pursue 3 "streams" of activity:

1. ***Apply models*** **to improve understanding of impacts of climate change/variability on hydrological cycle**
   - **Use California as the regional area of interest**
   - **Develop/Use quantitative metrics to evaluate skill**

2. ***Advance the science*** **by improving our models**
   - **improve subgrid scale physics (clouds, convection, precipitation)**
   - **Increase spatial resolution**
   - **Improve computational efficiency**

3. ***Make our results useful*** **to policymakers and water managers**

Methodology: High-Res Nested Models

Global Climate Model

LLNL

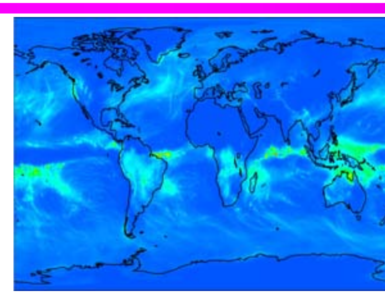

~75 km grid size

Climate models predict climate, including precipitation, and provide b.c. for downscaling.

Regional Climate Model

LLNL

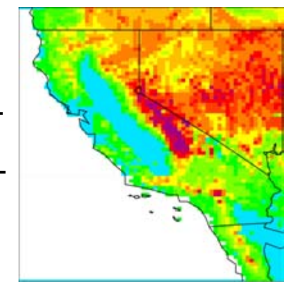

~20 km grid size

Hydrology model predicts runoff, given precipitation, etc. from climate model

Surface Hydrology Model

Partnering

~5 km grid size

Water Management Strategy Models

LLNL leads in high-res modeling

Higher resolution leads to improved simulated regional climate and hydrology

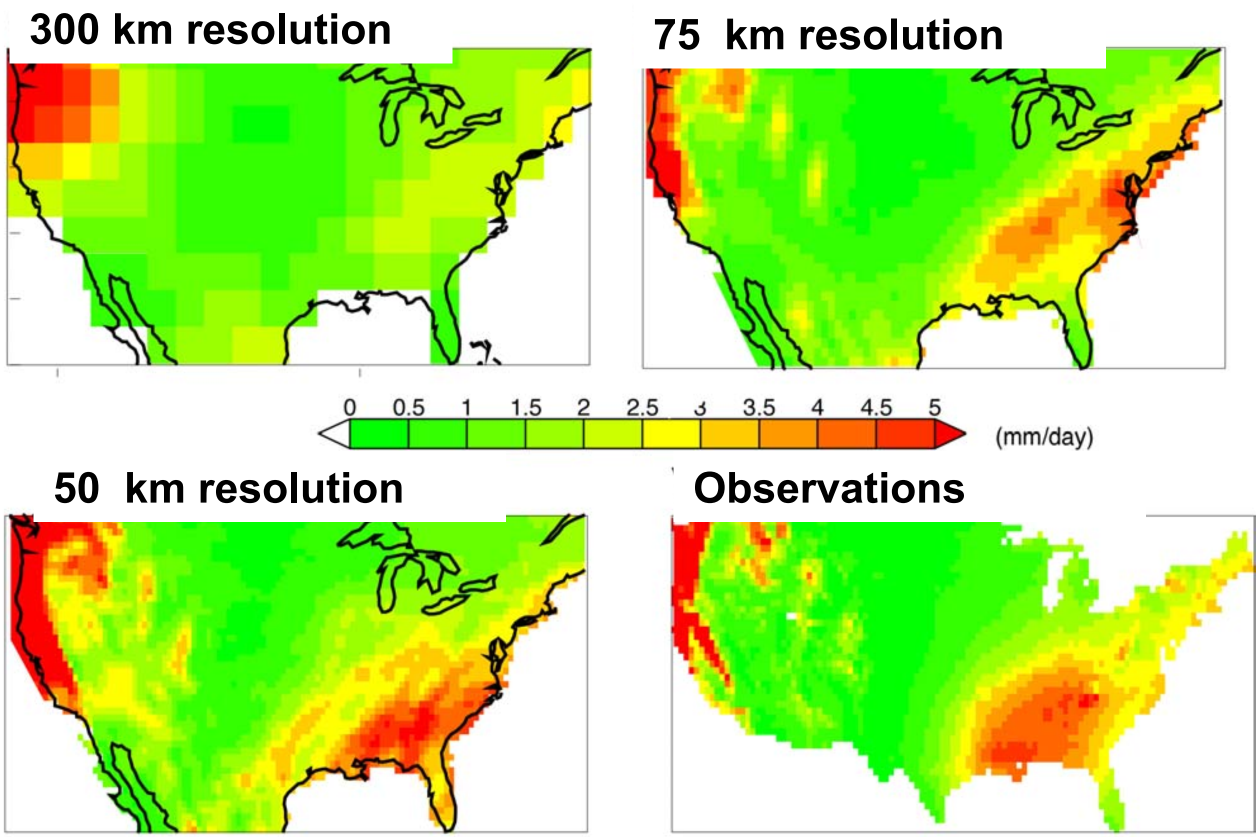

Simulated winter precipitation converges towards observations as model resolution becomes finer.

NASA / GSFC fvgcm climate code

- Finite-vol dynamical core
- Lagrangian vertical coord system
- Fast numerical algorithm
- Adaptable to massively parallel architectures
- “Highly scalable”

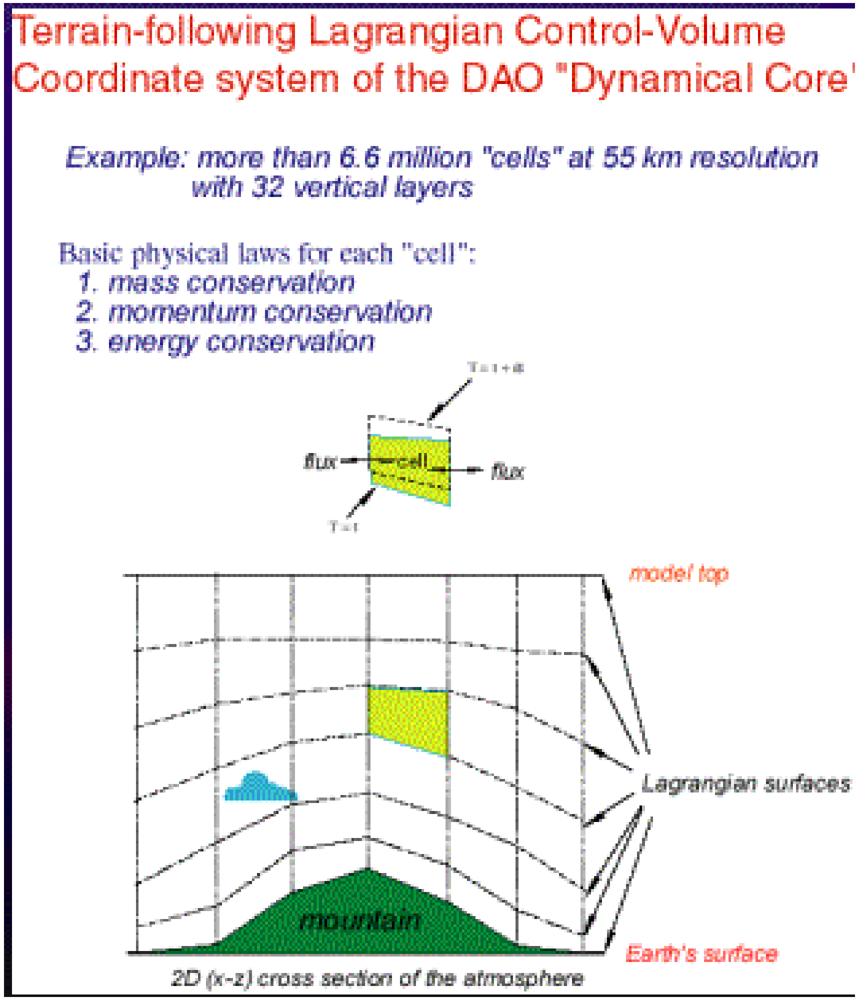

Climate Simulation Codes

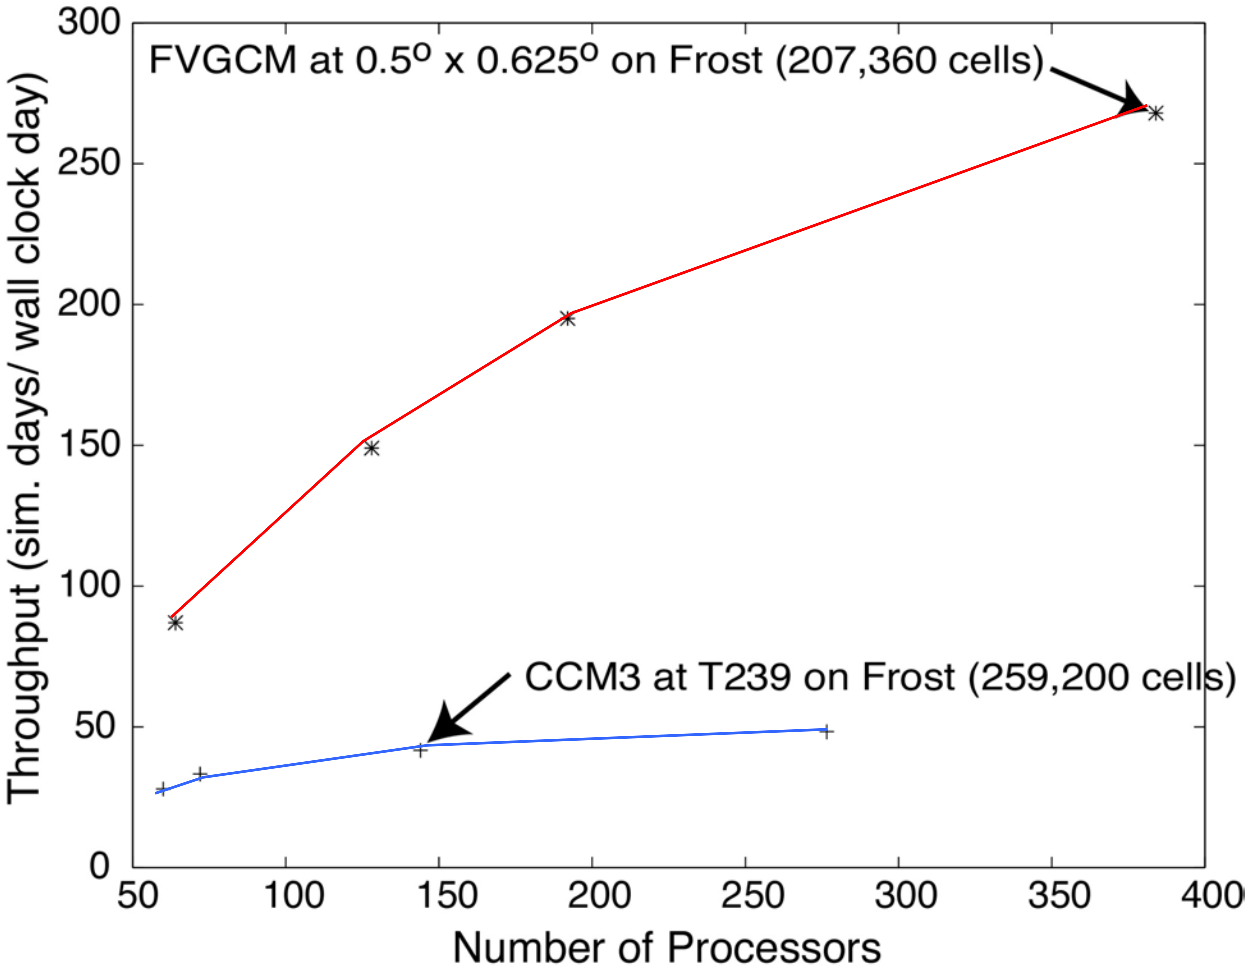

CCM3 = model used for previous high-resolution simulations
FVGCM = model we propose to use here

Computational Requirements

- 50 year climate simulations at ~ 60 km resolution
  - ~ 4 months wall time with 3 latitude stripes / node and 2 processors per node on MCR machine
  - 10 Terabytes data
- 50 years regional simulation @ 5 km resolution
- Surface hydrology simulation

Livermore MCR linux cluster

- 11.06 Tflops (1116 x 2 Intel 2.4 GHz processors)
- 4.6 TB memory
- 135 TB local disk
- 110 TB global disk (13 B:F)
- 320 MB/s MPI bandwidth and <5 μs latency over QsNet
- 120 MB/s transfers to Archive over dual Jumbo Frame Gb-Enet from each Login node
- 25 MB/s POSIX serial I/O to any file system
- Lustre file system with 4.48 GB/s delivered parallel I/O performance

Livermore MCR: Processing power

- Achieved 7.6 TFLOP/s (Linpack) on 12/3
- Dual 2.4Ghz Intel "Prestonia" with hyperthreading

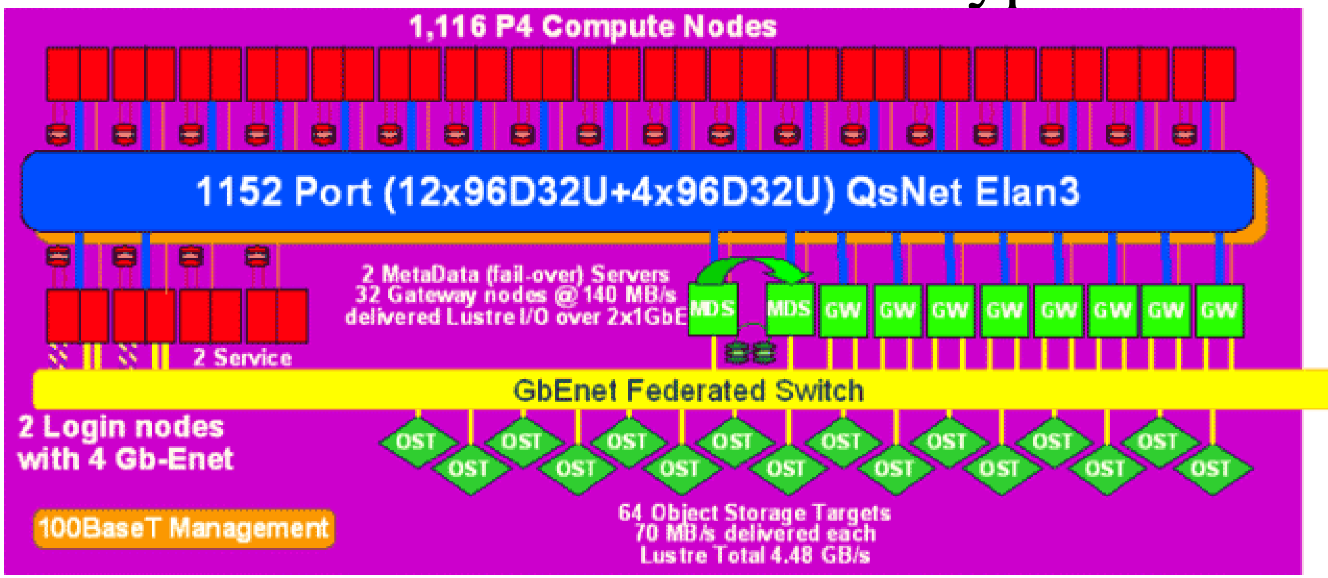

- Memory: 4GB PC-200 DDR SDRAM
- Local disk: 120GB ATA100 EIDE hard drive
- Quadrics QSNet Elan3 PCI adapter

Livermore MCR: I/O and Storage

Interim file system: 32x1.4 TB BlueArcs

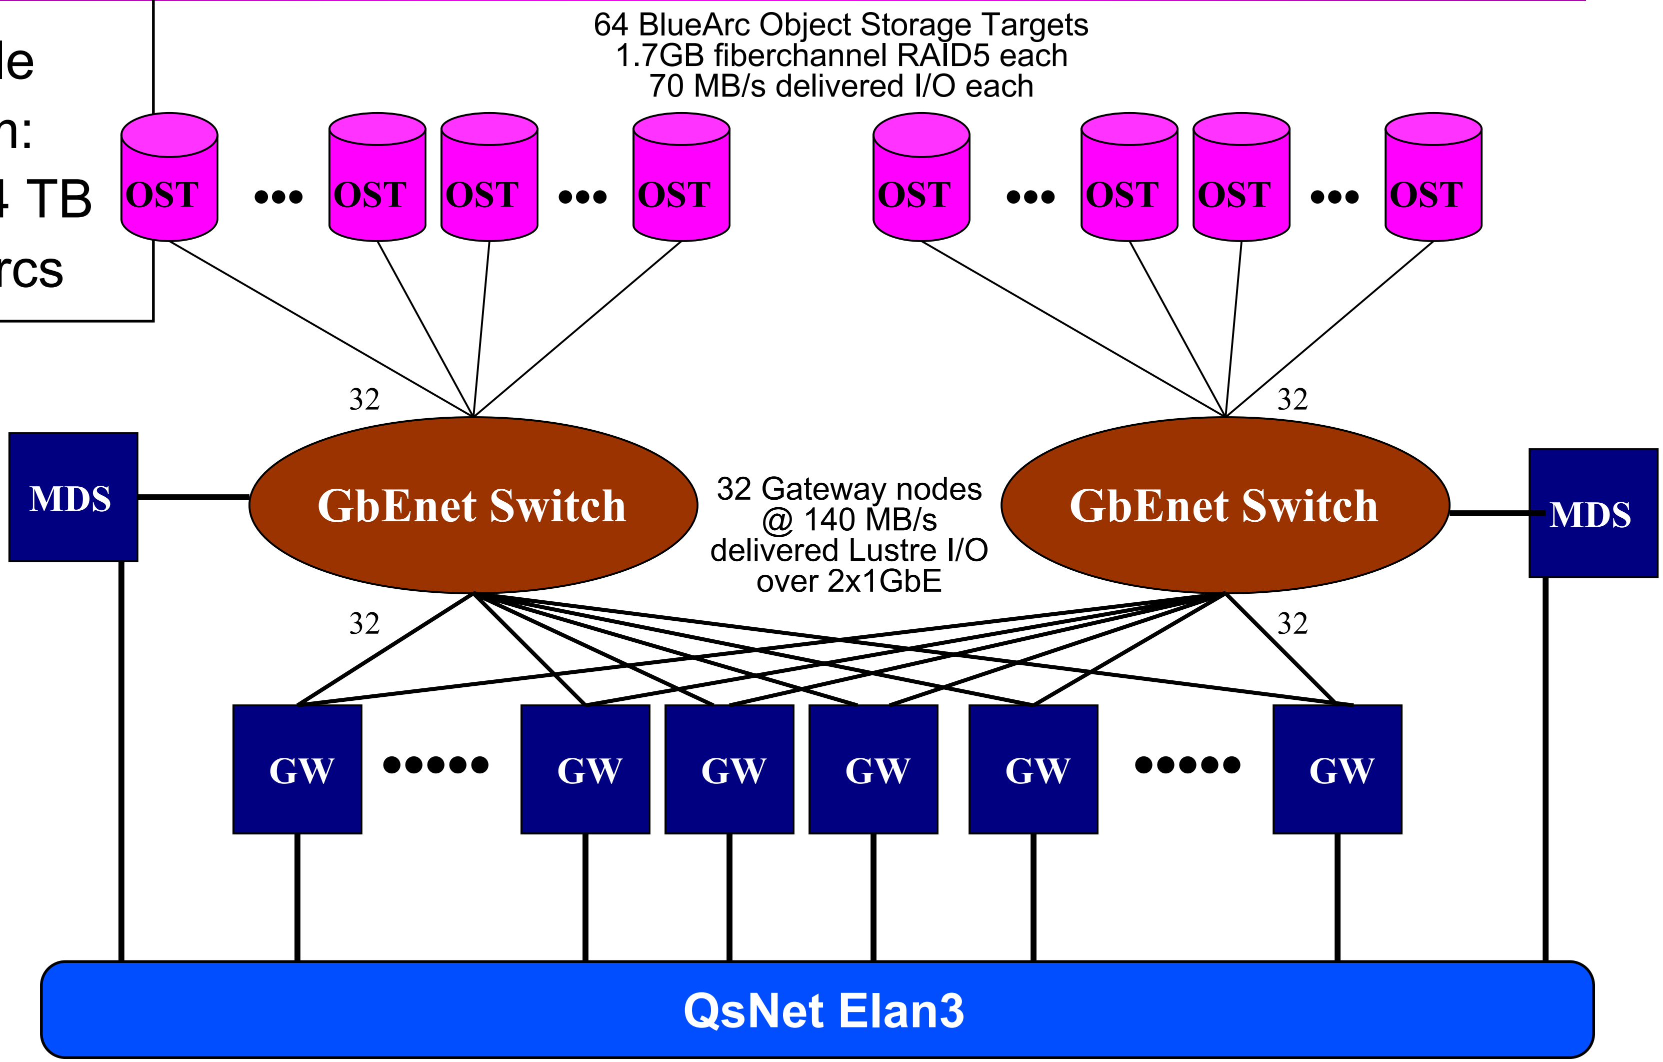

4.48 GB/s delivered parallel I/O performance!

Fvgcm on MCR: General Observations

- The Good:
  - MCR processors fast
  - fast fluid dynamics in Fvgcm code
  - tested at GSFC (SP, SGI, linux cluster) and LLNL (SP)
  - ~60 km resolution: 1 sim year / 3 days wall time w/ 120 proc
  - MPI2 (1-sided comm) appears to run ~ 20% faster than MPI1

- The Bad:
  - Porting took longer than expected (weeks, not hours)
  - Problems with MPI1 (2-sided comm) on mpi_gather
  - OpenMP thread problems: 1 cpu per node utilization
  - Fvgcm 1-d domain decomposition severely limits machine utilization

Fvgcm climate code mapped to MCR

- Domain decomposition:
  Mpi bands in latitude

  Example: “60 km” resolution

  576 lon x 360 lat x 32 vert

  (6,635,520 cells)

  120 shared mem nodes max

  OpenMP threads

  Mainly vertical

  2 proc / node on mcr cluster

  240 proc => 27648 cells/proc

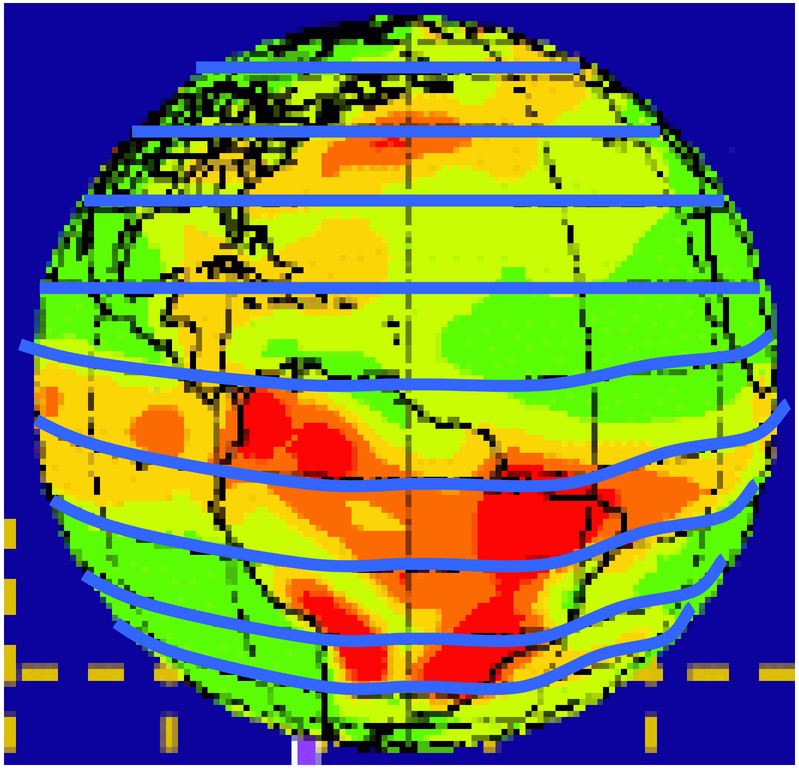

- Single processor output : global gather to master processor

Fvgcm on MCR

- Output through master only after global gather =>
  - buffer problems with mpi1
  - potentially slow output ?
- OpenMP threads fail (seg fault) with > 1 cpu / node
  - problem not yet solved
  - only observed on MCR machine

Scale-up Results

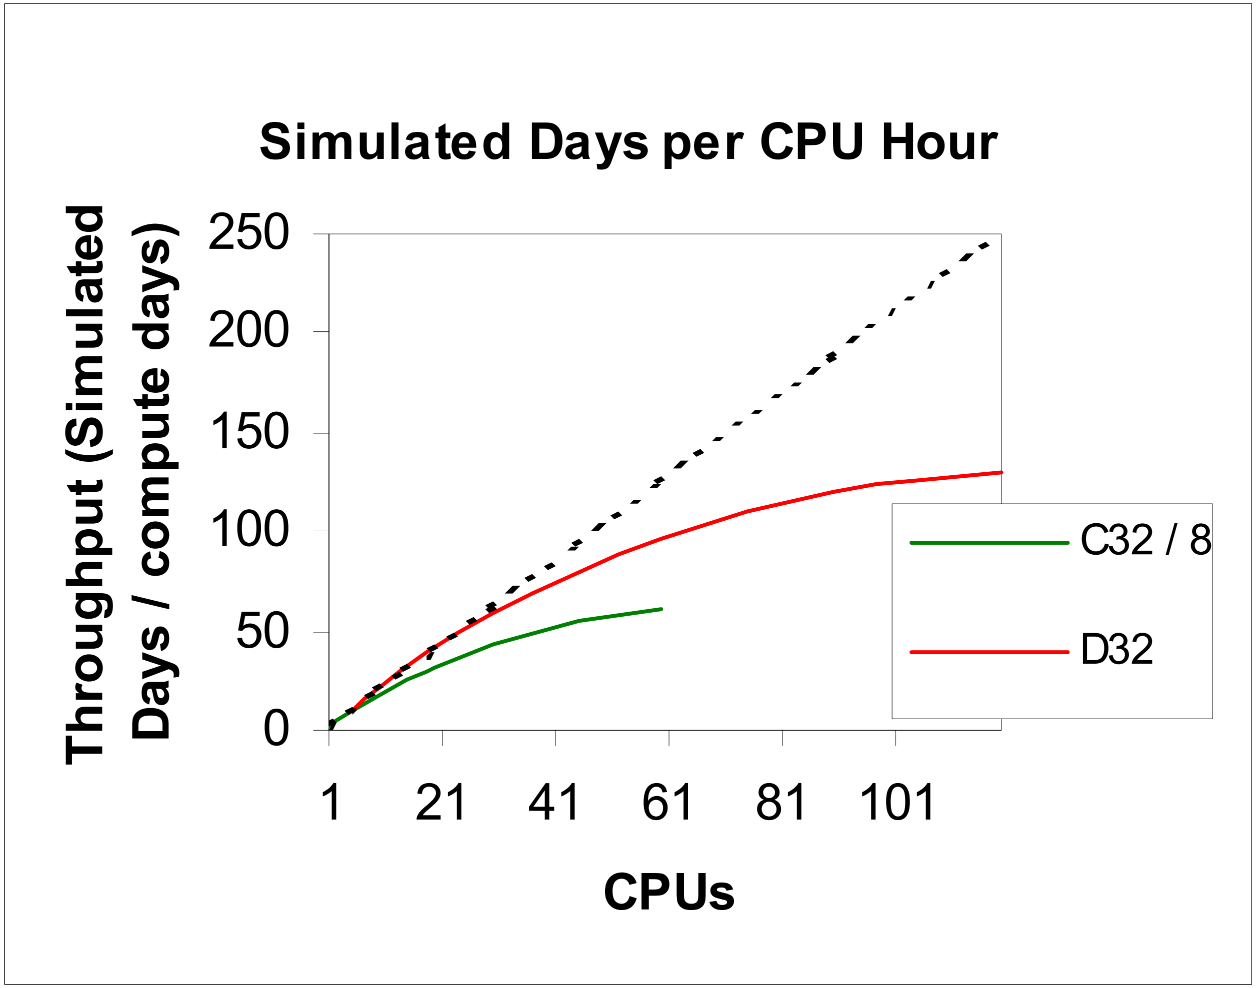

CPU and Storage Requirements

~ 60 km resolution: 576 x 361 x 32 (d32)

Simulated Climate Period	Wall Time	Data Storage Requirements (Terabytes)
1 month	5 hours	0.17
1 year	~ 3 days	2
50 years	~ 4 months	10

- ~2x Speedup with OpenMP ?

Initial Simulations: 60 km resolution (d32)

Mean large-scale winter precip rate (mm/day)

B32 resolution: ~ 240 km

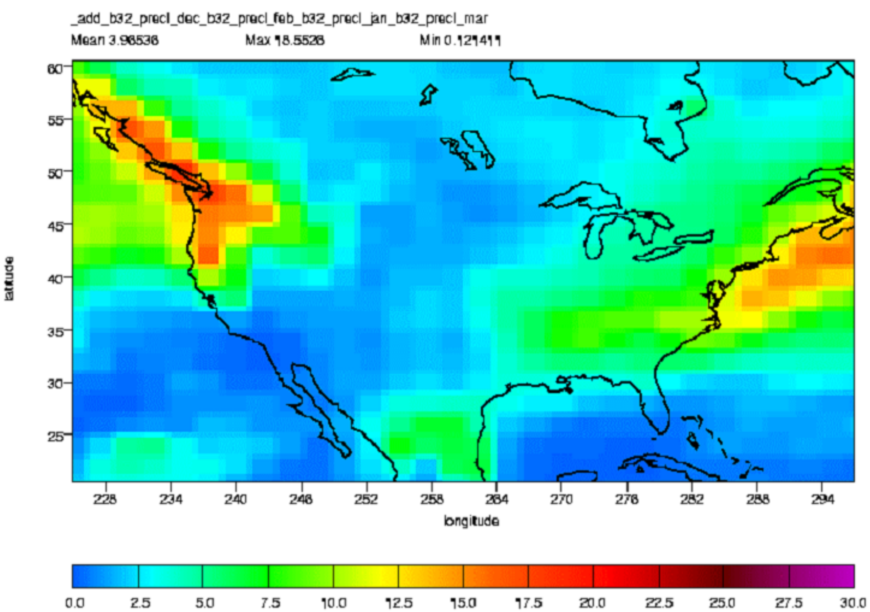

D32 resolution: ~ 60 km

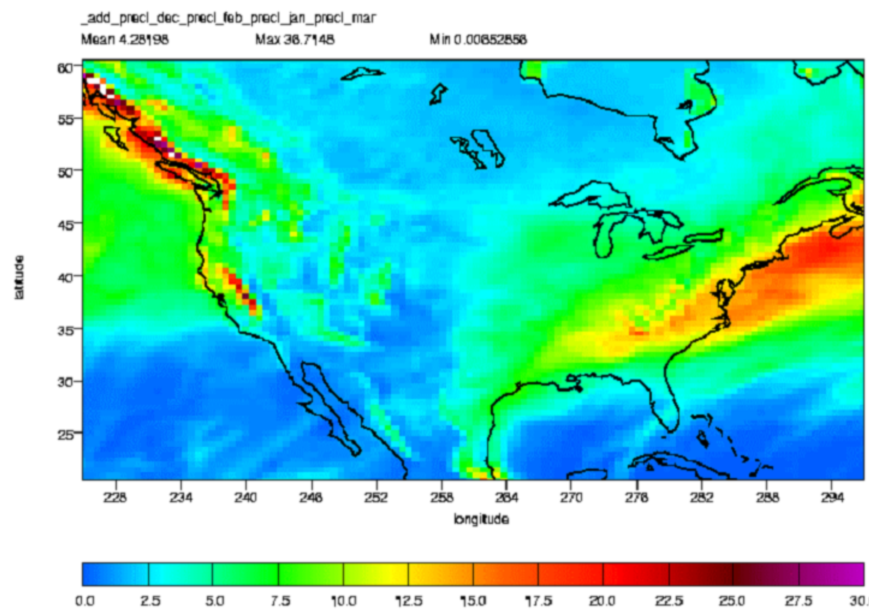

Summary

- Significant climate impact study: effects of increased CO_2 on California water supply
- Livermore MCR linux cluster has fast processors – good candidate for climate code
- Machine utilization limited by code domain decomp
- Unique MCR “features” cause problems with MPI and OpenMP; not seen on NASA/GSFC linux cluster
- Initial indications: total compute time quite good

Future Work

- California Water Project
  - Climate simulation for planning
  - May have high impact on agriculture and other industries in California
  - High-resolution simulation essential

- Long-term Climate Modeling
  - Better domain decomposition to take advantage of large node availability
  - Multi-resolution grids: regional impacts
  - I/O is important!
  - Move to commodity clusters a “good thing”: porting is still a major issue

UCRL-PRES-152874

Work performed under the auspices of the U. S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48